arXiv:1605.07906v1 [stat.ML] 25 May 2016

How priors of initial hyperparameters affect Gaussian process regression models

Zexun Chen*, Bo Wang

Department of Mathematics, University of Leicester, LE1 7RH, UK

Abstract

Gaussian Process Regression (GPR) is a kernel-based nonparametric method and has been proved to be effective and powerful. Its performance, however, relies on appropriate selection of kernel and the involving hyperparameters. The hyperparameters for a specified kernel are often estimated from the data via the maximum marginal likelihood. Unfortunately, the marginal likelihood functions are not usually convex with respect to the hyperparameters, therefore the optimization may not converge to global maxima. A common approach to tackle this issue is to use multiple starting points randomly selected from a specific prior distribution. Therefore, the choice of prior distribution may play a vital rule in the usefulness of this approach. In this paper, we study the sensitivity of prior distributions to the hyperparameter estimation and the performance of GPR. We consider different types of priors, including vague and data-dominated, for the initial values of hyperparameters for some commonly used kernels and investigate the influence of the priors on the performance of GPR models. The results show that the sensitivity of the hyperparameter estimation depends on the choice of kernels, but the priors have little influence on the performance of the GPR models in terms of predictability.

Keywords: Gaussian process regression, hyperparameters, kernel selection, prior distribution, maximum marginal likelihood.

1. Introduction

Over the last few decades, Gaussian Processes Regression (GPR) has been proven to be a powerful and effective method for non-linear regression problems due to many desirable properties, such as ease of obtaining and expressing uncertainty in predictions, the ability to capture a wide variety of behaviour through a simple parameterization, and a natural Bayesian interpretation [1]. Neal [2] revealed that many Bayesian regression models based on neural network converge to Gaussian Processes (GPs) in the limit of an infinite network.

*Corresponding author. Tel: +44-(0)77-6794-0465
Email address: zc74@le.ac.uk (Zexun Chen)

Therefore, GPs have been suggested as a replacement for supervised neural networks in non-linear regression [3, 4] and classification [3]. Furthermore, GPs have excellent capability for forecasting time series [5, 6].

However, GPR as a kernel-based nonparametric method, relies on appropriate selection of kernel [7] and the hyperparameters involved. Kernels contain our presumptions about the function we wish to learn and define the closeness and similarity between data points [8]. As a result, the choice of kernel has a profound impact on the performance of a GPR model, just as activation function, learning rate can affect the result of a neural network [7].

Once a kernel is selected, the unknown hyperparameters involved in the kernel need to be estimated from the training data. Although Monte Carlo methods can perform GPR without the need of estimating hyperparameters [4, 9, 6, 10], the common approach is to estimate the hyperparameters by means of maximum marginal likelihood [3] due to the high computational cost of Monte Carlo methods. Unfortunately marginal likelihood functions are not usually convex with respect to the hyperparameters, which means local optima may exist [11] and the optimized hyperparameters, which depend on the initial values, may not be the global optima [4, 6, 10, 12]. A common approach to tackle this issue is to use multiple starting points randomly selected from a specific prior distribution and after convergence choose the optimised values with the largest marginal likelihood as the estimates. Therefore, the choice of prior distribution may play a vital rule in the performance of GPR model. However, there exists little research in the literature to study the impact of the prior distributions on the hyperparameter estimation and the performance of GPR. Most researchers using GPR tend to choose a simple prior distribution based on their expert opinions and experiences, such as the Uniform distribution in the range of (0,1) [4, 6, 12].

In this paper, we study the sensitivity of prior distributions to the hyperparameter estimation and the performance of GPR. We consider different types of priors, including vague and data-dominated, for the initial values of hyperparameters for some commonly used kernels and investigate the influence of the priors on the performance of GPR model. The paper is organized as follows. Section 2 is a brief introduction to GPR. In Section 3, we discuss the key problem of the sensitivity of initial hyperparameters. Section 4 describes some different prior distributions for initial values, including both non-informative and data-dominated priors. Numerical experiments for different samples over different kernels are demonstrated and discussed in Section 5. Section 6 concludes the paper.

2. Background

2.1. Gaussian processes regression model

A Gaussian process is a collection of random variables, any finite number of which have (consistent) Gaussian distribution.

Mathematically, for any set S [1], a Gaussian process (GP) on S is a set of random variables $(f_x, x \in S)$ such that, for any $n \in \mathbb{N}$ and $x_1, \ldots, x_n \in S, (f_{x_1}, \ldots, f_{x_n})$ is (multivariate) Gaussian.

As a Gaussian distribution is specified by a mean vector and a covariance matrix, a GP is also fully determined by a mean function and a covariance function. In other words, we have:

Theorem 1 (Gaussian Processes)**.** *For any set S, any mean function $\mu : S \mapsto \mathbb{R}$ and any covariance function (also called kernel) $k : S \times S \mapsto \mathbb{R}$, there exists a GP $f(x)$ on S, s.t. $\mathbb{E}[f(x)] = \mu(x)$, $Cov(f(x_s), f(x_t)) = k(x_s, x_t), \forall x, x_s, x_t \in S$. It denotes $f \sim \mathcal{GP}(\mu, k)$.*

For a regression problem $y = f(x) + \varepsilon$, by Gaussian process method the unknown function f is assumed to follow a $\mathcal{GP}(\mu, k)$. Given n pairs of observations $(x_1, y_1), \ldots, (x_n, y_n)$, we have $\boldsymbol{y} = f(X) + \boldsymbol{\varepsilon}$, where $\boldsymbol{y} = [y_1, y_2, \ldots, y_n]^{\mathrm{T}}$ are the outputs, $X = [x_1, x_2, \ldots, x_n]^{\mathrm{T}}$ are the inputs, and $\boldsymbol{\varepsilon} = [\varepsilon_1, \varepsilon_2, \ldots, \varepsilon_n]^{\mathrm{T}}$ are independent identically distributed Gaussian noise with mean 0 and variance σ_n^2 [8]. It yields that the collection of functions $[f(x_1), \ldots, f(x_n)]$ follow a multivariate Gaussian distribution

$$[f(x_1), f(x_2), \ldots, f(x_n)]^{\mathrm{T}} \sim \mathcal{N}(\boldsymbol{\mu}, K),$$

where $\boldsymbol{\mu} = [\mu(x_1), \ldots, \mu(x_n)]^{\mathrm{T}}$ is the mean vector and K is the $n \times n$ covariance matrix of which the (i, j)-th element $K_{ij} = k(x_i, x_j)$.

To predict the function values $\boldsymbol{f}_* = [f_{*1}, \ldots, f_{*m}]^{\mathrm{T}}$ at the test locations $X_* = [x_{n+1}, \ldots, x_{n+m}]^{\mathrm{T}}$, the joint distribution of training observations $\boldsymbol{y}$ and predictive targets $\boldsymbol{f}_*$ are given by

$$\begin{bmatrix} \boldsymbol{y} \\ \boldsymbol{f}_* \end{bmatrix} \sim \mathcal{N}\left(\begin{bmatrix} \mu(X) \\ \mu(X_*) \end{bmatrix}, \begin{bmatrix} K(X, X) + \sigma_n^2 I & K(X_*, X)^{\mathrm{T}} \\ K(X_*, X) & K(X_*, X_*) \end{bmatrix} \right), \tag{1}$$

where $\mu(X) = \boldsymbol{\mu}$, $\mu(X_*) = [\mu(x_{n+1}), \ldots, \mu(x_{n+m})]^{\mathrm{T}}$, $K(X, X) = K$, $K(X_*, X)$ is an $m \times n$ matrix of which the (i, j)-th element $[K(X_*, X)]_{ij} = k(x_{n+i}, x_j)$, and $K(X_*, X_*)$ is an $m \times m$ matrix with the (i, j)-th element $[K(X_*, X_*)]_{ij} = k(x_{n+i}, x_{n+j})$. Thus the predictive distribution is

$$p(\boldsymbol{f}_* | X, \boldsymbol{y}, X_*) = \mathcal{N}(\hat{\mu}, \hat{\Sigma}), \tag{2}$$

$$\hat{\mu} = K(X_*, X)^{\mathrm{T}}(K(X, X) + \sigma_n^2 I)^{-1}(\boldsymbol{y} - \mu(X)), \tag{3}$$

$$\hat{\Sigma} = K(X_*, X_*) - K(X_*, X)^{\mathrm{T}}(K(X, X) + \sigma_n^2 I)^{-1} K(X_*, X). \tag{4}$$

In GPR method the mean function $\mu(x)$ is often assumed to be 0, then the predictive mean and variance can be given as

$$\hat{\mu} = K(X_*, X)^{\mathrm{T}}(K(X, X) + \sigma_n^2 I)^{-1} \boldsymbol{y}, \tag{5}$$

$$\hat{\Sigma} = K(X_*, X_*) - K(X_*, X)^{\mathrm{T}}(K(X, X) + \sigma_n^2 I)^{-1} K(X_*, X). \tag{6}$$

[1] Although S can be any set, it usually is $\mathbb{R}$ or $\mathbb{R}^n$. In this paper, we consider $S = \mathbb{R}$ only.

2.2. Kernels

From the view of Eq.(5) and Eq.(6), the kernel $k(\cdot,\cdot)$ plays a crucial role in the predictive mean and variance. As discussed in [8], kernels contain our presumptions about the function we wish to learn and define the closeness and similarity between data points. As a result, the choice of kernel has a profound impact on the performance of a GPR model, just as activation function, learning rate can affect the result of a neural network [7]. Some commonly used kernels are listed as follows.

2.2.1. Squared exponential

The most widely-used kernel in GPR is Squared Exponential (SE), which is defined as

$$k_{SE}(x,x') = s_f^2 \exp(-\frac{(x-x')^2}{2\ell^2}),$$

where s_f is the signal variance and can also be considered as an output-scale amplitude [13] and the parameter ℓ is the input (length or time) scale [13].

2.2.2. Periodic

Periodic kernel (PER) is used to model functions which exhibit a periodic pattern. It is given by

$$k_{PER}(x,x') = s_f^2 \exp(-\frac{2\sin^2(\pi\frac{(x-x')}{p})}{\ell^2}),$$

where p is the period of the function and the parameters s_f and ℓ have the same meaning as in SE.

2.2.3. Local periodic

As seen in [11], positive semi-definite kernels are closed under addition and multiplication. Local Periodic (LP) is such a composite kernel which is obtained by multiplying SE and PER [11]. That is,

$$k_{LP}(x,x') = k_{SE}(x,x') \times k_{PER}(x,x').$$

It is a well-known kernel to capture locally periodic structure of data hence can be applied to many kernel-based models.

2.2.4. Spectral mixture

The Spectral Mixture (SM) kernel was introduced by Wilson [12] and is defined as a scaled mixture of Q Gaussians:

$$k_{SM}(x,x') = \sum_{q=1}^{Q} w_q \exp(-2\pi^2(x-x')^2\nu_q)\cos(2\pi(x-x')\mu_q),$$

where w_q's are the weights, the inverse means $1/\mu_q$ represent component period and each inverse standard deviation $1/\sqrt{\nu_q}$ represents a length scale [12].

2.3. Estimation of hyperparameters

In GPR models, the hyperparameters involved in the kernel need to be estimated from the training data. Although Monte Carlo methods can perform GPR without the need of estimating hyperparameters [4, 9, 6, 10], the common approach is to estimate them by means of maximum marginal likelihood due to the high computational cost of Monte Carlo methods.

Following the GP assumption, the distribution of the training outputs is given as

$$p(\boldsymbol{y}|X,\boldsymbol{\theta}) = \mathcal{N}(\mathbf{0}, \Sigma_\theta), \tag{7}$$

where $\Sigma_\theta = K + \sigma_n^2 I$ and $\boldsymbol{\theta}$ is the collection of the unknown hyperparameters. Therefore, the negative log marginal likelihood (nlml) is

$$\mathcal{L}(\boldsymbol{\theta}) = -\log p(\boldsymbol{y}|X,\boldsymbol{\theta}) = \frac{1}{2}\boldsymbol{y}^{\mathrm{T}}\Sigma_\theta^{-1}\boldsymbol{y} + \frac{1}{2}\log\det\Sigma_\theta + \frac{n}{2}\log 2\pi, \tag{8}$$

and the partial derivatives of nlml with respect to the hyperparameters are given by

$$\frac{\partial}{\partial\theta_i}\mathcal{L}(\boldsymbol{\theta}) = \frac{1}{2}\mathrm{tr}(\Sigma_\theta^{-1}\frac{\partial\Sigma_\theta}{\partial\theta_i}) - \frac{1}{2}\boldsymbol{y}^{\mathrm{T}}\Sigma_\theta^{-1}\frac{\partial\Sigma_\theta}{\partial\theta_i}\Sigma_\theta^{-1}\boldsymbol{y}. \tag{9}$$

3. Sensitivity of prior distributions for initial hyperparameters

For many kernels the likelihood function (8) is not convex with respect to the hyperparameters, therefore the optimisation algorithm may converge to a local optimum whereas the global one may provide better results [11]. As a result the optimised hyperparameters achieved by maximum likelihood estimation and the performance of GPR may depend on the initial values of the optimisation algorithm [4, 6, 10, 12].

A common strategy adopted by most GPR practitioners is a heuristic method. That is, the optimisation is repeated using several initial values generated randomly from a simple prior distribution, which is often selected based on their expert opinions and experiences. The final estimates of the hyperparameters are the ones with the largest likelihood values after convergence [4, 6, 12]. It is therefore interesting to know how prior distributions affect the performance of GPR since the above strategy can not guarantee a global maximum of the likelihood function is found, or the sensitivity of prior distributions to the performance of GPR, which, to the best of our knowledge, has not been studied in the literature. In this paper, we consider several different priors and study their influences to the estimates of the hyperparameters and the performance of GPR models for some commonly used kernels.

4. Prior distributions of initial hyperparameters

The prior distributions considered include non-informative [7] and data-dominated [12], which are briefly introduced as follows.

4.1. Vague priors

In the cases when there is little information about the data, vague prior distributions are often selected with the intention that they should have slight or no influence on the inferences [14, 15]. Many justifications and interpretations of non-informative priors have been proposed over the years, including invariance [16] and maximum entropy [17]. However, with small amount of data, the use of non-informative prior may be problematic and a vague prior distribution may lead to significant influence on any inference made because the results are easily sensitive to the selection of prior distributions [14].

Let θ_i be a generic notation for a hyperparameter in a given kernel. Below list the weakly-informative prior distributions which will be discussed in our study.

Prior 1

$$\theta_i \sim \text{Uniform}(0, 1).$$

This is probably the most common prior distribution. Actually, it is not strictly a 'vague' prior since the range of the distribution is restricted. However, this prior is very widely-used for the estimation of the unknown parameters in the GPR models.

Prior 2

$$\log(\theta_i) \sim \text{Uniform}(-1, 1).$$

This prior distribution is uniform on the log hyperparameters in $(-1, 1)$, so the range of the hyperparameters is $(1/\mathrm{e}, \mathrm{e})$.

Prior 3

$$\log(\theta_i) \sim \text{Uniform}(-10, 10).$$

This prior is similar to Prior 2 but has much larger range. So the range of the hyperparameters is approximately $(0, \mathrm{e}^{10})$.

Prior 4

$$\theta_i \sim \mathcal{N}(0, 1).$$

The standard normal prior is also a popular and simple choice. It is not strictly a 'vague' prior either, and cannot be used for positive parameters.

Prior 5

$$\frac{\pi}{\theta_i} \sim \text{Uniform}(0, 1).$$

This prior is specified for the period parameter for kernels that contain periodic part. The range of the parameter is $(\pi, +\infty)$.

Prior 6

$$\log(\frac{\pi}{\theta_i}) \sim \text{Uniform}(-5, 5).$$

This prior is also specified for the period parameter. It is similar to Prior 5 but with a range $(\pi\mathrm{e}^{-5}, \pi\mathrm{e}^{5})$.

4.2. Data-dominated priors

Data-dominated priors are incorporated with some information inferred from training data, such as the possible range of the initial hyperparameters. The following data-dominated priors will be used in this study.

Prior 7

$$\theta_i \sim \text{Uniform}(0, \text{Nyq}).$$

This prior is also specified for the period parameter and is based on Nyquist frequency [18], where Nyq equals half the sampling rate of the data, or half the largest interval between input points if the data are not regularly sampled [12]. Nyquist frequency can be used to find the approximate period of data in signal processing and spectral analysis. For example, Wilson [12] used this prior to initialise SM kernel.

Prior 8

$$\frac{1}{\theta_i} \sim \mathcal{TN}(\text{MaxI}),$$

where $\mathcal{TN}(\text{MaxI})$ is the truncated normal distribution with mean proportional to the maximal range of the inputs (MaxI) [12]. It is an improved version of Prior 4 and is used by Wilson [12] for the length scale in SM kernel.

Prior 9

$$\frac{\pi}{\theta_i} \sim \text{Uniform}(\frac{\pi}{\text{MaxI}}, \pi\text{Nyq}).$$

This prior is also specified for the period parameters and has the range $(1/\text{Nyq}, \text{MaxI})$. It was first used in [12] to find a suitable range for the initial hyperparameters.

5. Experiments

5.1. Experiments using samples from Gaussian processes

In this section, we study how the priors of initial hyperparameters affect the estimates of the hyperparameters and the performance of GPR models using data generated from specified Gaussian processes. Since the true models are known, the accuracy of the estimates can be compared.

Letting $x_i = i$ for $i = 1, 2, \ldots, 400$, we generate samples $\{y_i\}$ from GPs with zero mean and SE and PER kernels, respectively. These two kernels are used as demonstration because SE is the most widely-used kernel in GPR while PER is the simplest kernel which may suffer from the problem of local optima in optimisation procedure because integer multiples of the true period, such as harmonics, are often local optima [11].

To evaluate the influences of the prior distributions on the hyperparameter estimation, ten values randomly generated from each prior distribution discussed in Section 4 (where applicable) are used as the starting values for the maximum likelihood procedure, implemented by Conjugate Gradient algorithm. Among the ten estimates after the procedure is converged the one with the largest maximum likelihood is chosen as the optimal estimate, denoted by $\boldsymbol{\theta}_{final}$, and is compared with $\boldsymbol{\theta}_{act}$.

To study the impact of the priors on the predictability of GPR, we consider two types of prediction: interpolation and extrapolation. Denote the whole data set by $\Omega = \{(i, y_i); i = 1, 2, \ldots, 400\}$. For interpolation, the test set is given by $\mathcal{D}_{I2} = \{(i, y_i); i = 5j+1, j = 0, 1, \ldots, 79\}$ and the training set is $\mathcal{D}_{I1} = \Omega - \mathcal{D}_{I2}$. For extrapolation, the training set is $\mathcal{D}_{E1} = \{(i, y_i); i = 1, 2, \ldots, 320\}$ and the test set is $\mathcal{D}_{E2} = \Omega - \mathcal{D}_{E1}$.

The predicted values are then compared with the actual values. There are several ways to evaluate the accuracy of the predictions. The simplest one is the root mean squared error (RMSE), which is defined as

$$\text{RMSE} = \sqrt{\frac{1}{m}\sum_{i=1}^{m}(\hat{y}_i - y_i)^2},$$

where $\{\hat{y}_i\}$ and $\{y_i\}, i = 1, 2, \ldots, m$, are the predicted mean values and the actual test values respectively. However, the RMSE can be affected seriously by the overall scale of the output values, so we utilize the standardized root mean squared error (SRMSE) which is normalized by the standard deviation of $\{y_i\}$, i.e.

$$\text{SRMSE} = \frac{\text{RMSE}}{\sigma_y},$$

where σ_y is the standard deviation of $\{y_i\}, i = 1, 2, \ldots, m$. This implies any model which can provide the prediction close to the sample mean of the test targets to have a SRMSE of approximately 1 [8]. In other words, any prediction model with the SRMSE around 1 is satisfactory.

Another measure which can take account of both predictive mean and predictive variance is log loss. As the predictive distribution for each test point is Gaussian, its log loss is defined as

$$\text{LL} = \frac{1}{2}\log(2\pi\hat{\sigma}_i^2) + \frac{(y_i - \hat{y}_i)^2}{2\hat{\sigma}_i^2},\ i = 1, 2, \ldots, m,$$

where $\{\hat{\sigma}_i\}, i = 1, 2, \ldots, m$ are the predictive variances. This loss can be standardized by subtracting the loss that could be obtained by the null model which predicts using a Gaussian with the sample mean and sample variance of the training outputs [8]. And the mean standardized log loss (MSLL) is the average of the standardized log loss for $i = 1, 2, \ldots, m$. Therefore, the MSLL is zero for null model, and the smaller it is the better a model is in terms of loss [8].

5.1.1. Squared Exponential kernel

As can be seen in Section 4, not all of the priors are suitable for every hyperparameter. Therefore for SE kernel, we use **Prior** 1, **Prior** 2, **Prior** 3 for both hyperparameters $[\ell, s_f]$. The data are generated using $\boldsymbol{\theta}_{act} = [\ell, s_f] = [5, 2]$.

To compare $\boldsymbol{\theta}_{act}$ and $\boldsymbol{\theta}_{final}$, Figure 1 illustrates their visual positions, where "$\square$" represents $\boldsymbol{\theta}_{act}$, "$\star$" represents $\boldsymbol{\theta}_{final}$, the "$+$"s are the intermediate

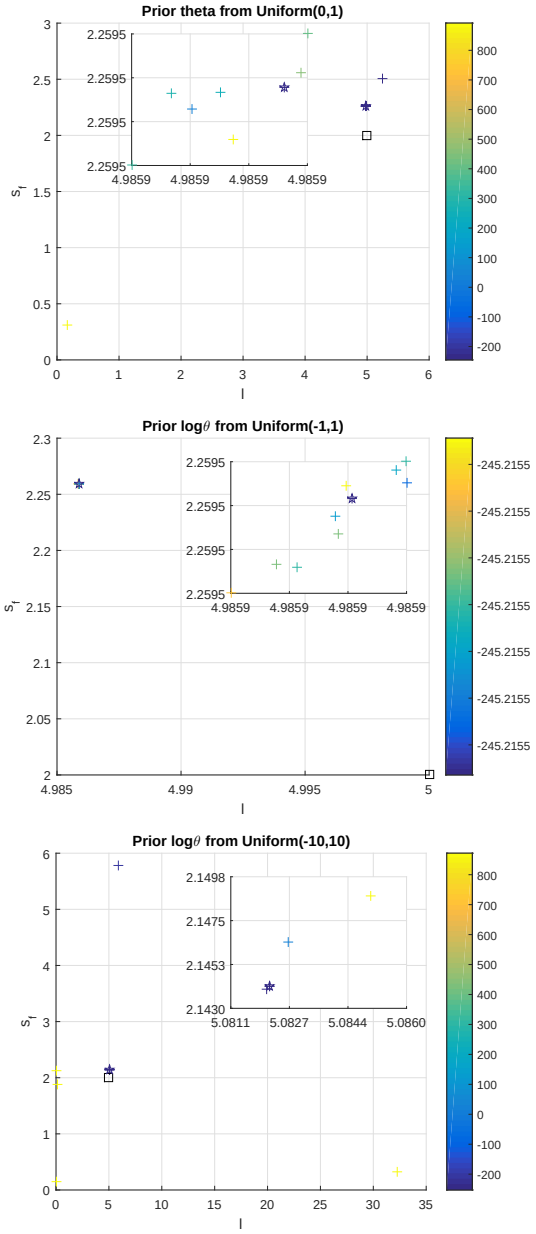


Figure 1: Positions of the estimated hyperparameters for the SE kernel. Top to bottom: Priors 1, 2 and 3.

values during the process of optimization and the color of the symbols stands for the value of the negative log marginal likelihood (nlml).

Apparently, regardless of the priors, the optimisation converges very fast and the estimated hyperparameter $\boldsymbol{\theta}_{final}$ is always very close to $\boldsymbol{\theta}_{act}$.

We now test the prediction performance by the GPR with SE kernel. Only **Prior 1** is used since the estimated hyperparameters from different priors are almost the same. The samples are generated using two GP models with two different hyperparameters: $\boldsymbol{\theta}_{act} = [5, 2]$ and $\boldsymbol{\theta}_{act} = [15, 7]$, respectively. And the above experiment is repeated 20 times and the average results are reported in Table 1. As demonstration, Figure 2 shows typical predictions of GP and corresponding SRMSE and MSLL.

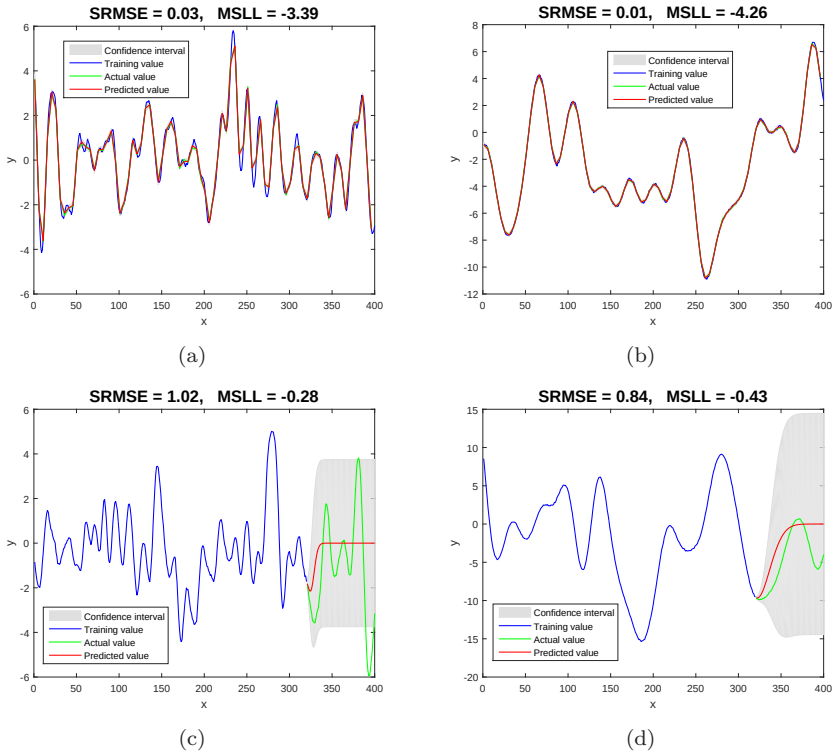


Figure 2: GP predictions with SE kernel using Prior 1. (a) and (b): interpolations; (c) and (d): extrapolations. (a) and (c): $\boldsymbol{\theta}_{act} = [5, 2]$; (b) and (d): $\boldsymbol{\theta}_{act} = [15, 7]$.

It is obvious that the mean estimate of $\boldsymbol{\theta}_{final}$ is very close to $\boldsymbol{\theta}_{act}$ with small standard errors for both cases, and the GPR model performs well and stably for both interpolation and extrapolation predictions.

Table 1: Results of GP predictions with SE kernel (the standard errors are given in the brackets)

Interpolation				
$\boldsymbol{\theta}_{act}$		$\boldsymbol{\theta}_{final}$	SRMSE	MSLL
ℓ	5	4.97 (0.133)	0.03 (0.004)	-3.48 (0.130)
s_f	2	2.00 (0.187)		
ℓ	15	14.95 (0.467)	0.01 (0.002)	-4.74 (0.229)
s_f	7	6.94 (1.110)		
Extrapolation				
$\boldsymbol{\theta}_{act}$		$\boldsymbol{\theta}_{final}$	SRMSE	MSLL
ℓ	5	4.98 (0.165)	1.02 (0.141)	-0.14 (0.114)
s_f	2	1.97 (0.210)		
ℓ	15	15.01 (0.495)	1.19 (0.535)	-0.56 (0.322)
s_f	7	7.01 (1.242)		

5.1.2. Periodic Kernel

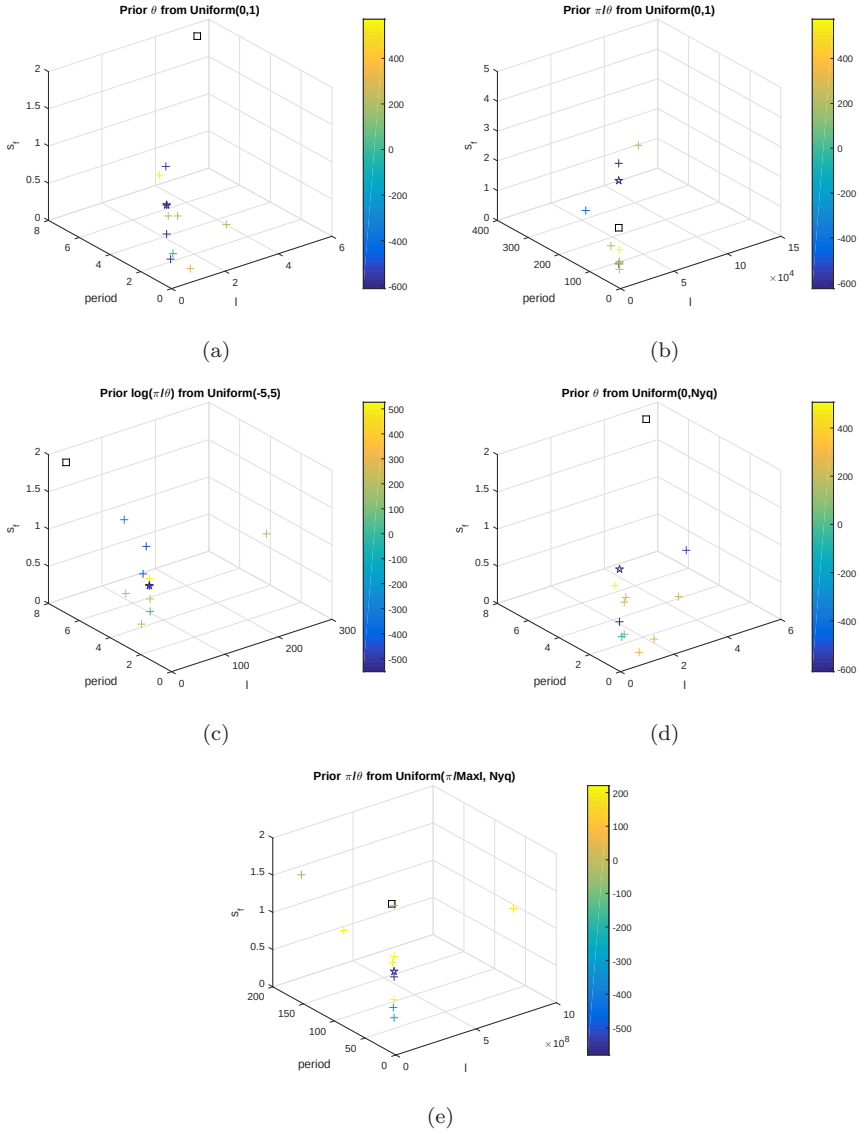
Three parameters $[\ell, period, s_f]$ are involved in the PER kernel. We consider five priors (**Prior** 1, **Prior** 5, **Prior** 6, **Prior** 7 and **Prior** 9) for the *period* term and **Prior 1** for the parameters ℓ and s_f. In the following experiment, the data are generated using the true parameters $\boldsymbol{\theta}_{act} = [5, 7, 2]$.

Figure 3 shows the visual positions of $\boldsymbol{\theta}_{act}$ and $\boldsymbol{\theta}_{final}$, where the symbols have the same meanings as in Figure 1. It can be seen that, for all the priors considered, the estimates $\boldsymbol{\theta}_{final}$ are always far away from the true value $\boldsymbol{\theta}_{act}$. Therefore, it is difficult to achieve the global maximum by the maximum marginal likelihood method for the PER kernel, and the estimates are very sensitive to prior distributions of the initial hyperparameters.

The same strategy as for the SE kernel is used to test the prediction performance by the GPR with PER kernel, and the results are reported in Table 2. It can be seen that, consistent with the above findings, the averages of the estimated hyperparameters are very different than the true values, which confirms that the estimates obtained by numerical optimization of likelihood function are biased. However, both the means and standard deviations of SRMSE and MSLL are very small, which indicates that the GPR models perform very well and stably for both interpolation and extrapolation, despite the poor estimates of the hyperparameters. Therefore, although the parameter estimation for the PER kernel is sensitive to prior distributions, the GPRs still provide good results and the performance is hardly influenced by the choice of priors.

5.2. Experiments using samples from time series

It is of interest to investigate how prior distributions of the hyperparameters influence the predictability of GPR if the data are generated from other models.



(a) (b) (c) (d) (e)

Figure 3: Positions of the estimated hyperparameters for the PER kernel. The priors for the *period* term are: (a) **Prior** 1, (b) **Prior** 5, (c) **Prior** 6, (d) **Prior** 7 and (e) **Prior** 9.

Table 2: Results of GP predictions with PER kernel (the standard errors are given in the brackets)

Interpolation					
Prior	$\boldsymbol{\theta}_{act}$		$\boldsymbol{\theta}_{final}$	SRMSE	MSLL
Prior 1	ℓ	5	0.25 (0.170)	0.35 (0.454)	-1.44 (1.205)
	period	7	1.98 (2.807)		
	s_f	2	2.41 (2.453)		
Prior 5	ℓ	5	4.24 (8.968)	0.48 (0.815)	-1.63 (1.365)
	period	7	4.81 (1.293)		
	s_f	2	95.98 (205.308)		
Prior 6	ℓ	5	1.19 (0.713)	0.28 (0.252)	-1.50 (0.705)
	period	7	2.98 (2.288)		
	s_f	2	3.58 (6.246)		
Prior 7	ℓ	5	1.45 (1.289)	0.28 (0.253)	-1.48 (0.696)
	period	7	0.34 (0.143)		
	s_f	2	1.51 (0.838)		
Prior 9	ℓ	5	1.67 (1.978)	0.28 (0.252)	-1.51 (0.712)
	period	7	13.54 (16.614)		
	s_f	2	39.26 (80.484)		
Extrapolation					
Prior	$\boldsymbol{\theta}_{act}$		$\boldsymbol{\theta}_{final}$	SRMSE	MSLL
Prior 1	ℓ	5	1.23 (1.048)	0.14 (0.041)	-1.98 (0.287)
	period	7	0.40 (0.254)		
	s_f	2	1.43 (1.230)		
Prior 5	ℓ	5	17.87 (47.950)	0.24 (0.119)	-1.52 (0.492)
	period	7	7.73 (3.126)		
	s_f	2	56.50 (214.170)		
Prior 6	ℓ	5	2.01 (2.023)	0.24 (0.120)	-1.51 (0.493)
	period	7	2.20 (1.323)		
	s_f	2	2.38 (3.984)		
Prior 7	ℓ	5	24.18 (87.495)	0.17 (0.103)	-1.91 (0.538)
	period	7	0.25 (0.117)		
	s_f	2	139.47 (612.532)		
Prior 9	ℓ	5	4.07 (6.696)	0.18 (0.106)	-1.82 (0.563)
	period	7	5.80 (2.898)		
	s_f	2	4.71 (9.551)		

We consider a simple time series model ARMA(2,1) with autoregressive coefficient $[0.8, -0.45]$ and moving average coefficient -0.5, and generate 400 samples $\{y_i, i = 1, 2, \ldots, 400\}$ with $x_i = i$ and the starting values $y_1 = y_2 = 1$. Here we consider extrapolation only as this type of prediction is more meaningful in time series modelling. We select the first 320 data points as the training data and the rest as the test data. The GPR models are applied using two composite kernels: local periodic (LP) and spectral mixture (SM) with 4 components, both of which are known as useful kernels for data with complex pattern [12]. For LP kernel, different priors are used for the *period* parameter while Prior 1 is used for all the remaining parameters. For SM kernel, as shown in Section 2, three parameters $[w_q, \mu_q, \nu_q]$ are involved. However, w_q can be initialised as constants proportional to the standard deviation of the data [12]. Therefore we only focus on the remaining hyperparameters μ_q and ν_q. We denote

$$[\mu_q, \sqrt{\nu_q}] \sim \begin{cases} \text{PS51}, & \text{if } \mu_q \text{ is from } \textbf{Prior } 5, \sqrt{\nu_q} \text{ is from } \textbf{Prior } 1 \\ \text{PS61}, & \text{if } \mu_q \text{ is from } \textbf{Prior } 6, \sqrt{\nu_q} \text{ is from } \textbf{Prior } 1 \\ \text{PS71}, & \text{if } \mu_q \text{ is from } \textbf{Prior } 7, \sqrt{\nu_q} \text{ is from } \textbf{Prior } 1 \\ \text{PS91}, & \text{if } \mu_q \text{ is from } \textbf{Prior } 9, \sqrt{\nu_q} \text{ is from } \textbf{Prior } 1 \\ \text{PS58}, & \text{if } \mu_q \text{ is from } \textbf{Prior } 5, \sqrt{\nu_q} \text{ is from } \textbf{Prior } 8 \\ \text{PS68}, & \text{if } \mu_q \text{ is from } \textbf{Prior } 6, \sqrt{\nu_q} \text{ is from } \textbf{Prior } 8 \\ \text{PS78}, & \text{if } \mu_q \text{ is from } \textbf{Prior } 7, \sqrt{\nu_q} \text{ is from } \textbf{Prior } 8 \\ \text{PS98}, & \text{if } \mu_q \text{ is from } \textbf{Prior } 9, \sqrt{\nu_q} \text{ is from } \textbf{Prior } 8 \end{cases}$$

where PS78 is the priors used by Wilson [12].

For comparison of the performance, the prediction is also performed using the true model ARMA(2,1) with the true parameters. The experiment is repeated 20 times and the averages and the standard deviations are reported in Tables 3 and 4, respectively.

Table 3: Results of GP predictions with LP kernel for ARMA data (the standard errors are given in the brackets)

	GPR with LP kernel		ARMA(2,1)	
Priors	SRMSE	MSLL	SRMSE	MSLL
Prior 1	1.006 (0.0221)	-0.001 (0.0151)		
Prior 5	1.007 (0.0223)	-0.001 (0.0154)		
Prior 6	1.006 (0.0219)	-0.001 (0.0150)	1.006 (0.0143)	-0.002(0.0077)
Prior 7	1.006 (0.0219)	-0.001 (0.0150)		
Prior 9	1.005 (0.0225)	-0.002 (0.0149)		

The results show that for both LP and SM kernels, the performance of the GPR models has no significant differences using different prior distributions, and is comparable to that by the true model. In other words, the performance of GPR models is not sensitive to the choice of prior distributions and is as good as the true model as far as this experiment concerns.

Table 4: Results of GP predictions with SM kernel for ARMA data (the standard errors are given in the brackets)

	GPR with SM kernel		ARMA(2,1)	
Priors	SRMSE	MSLL	SRMSE	MSLL
PS51	1.008 (0.0251)	0.001 (0.0207)	1.006 (0.0143)	-0.002 (0.0077)
PS61	1.007 (0.0236)	-0.001 (0.0187)		
PS71	1.009 (0.0255)	0.001 (0.0210)		
PS91	1.006 (0.0252)	-0.002 (0.0188)		
PS58	1.036 (0.0498)	0.038 (0.0415)		
PS68	1.043 (0.0509)	0.036 (0.0450)		
PS78	1.019 (0.0350)	0.012 (0.0351)		
PS98	1.032 (0.0490)	0.028 (0.0441)		

6. Conclusion

In this paper, we conducted the simulation studies to investigate the influences of various prior distributions of the initial hyperparameters in GPR models to the parameter estimation and the predictability of the models when numerical optimisation of likelihood function was utilised. Nine commonly used priors and four kernels, including two basic kernels (SE and PER) and two complicated kernels (LP and SM), were considered.

The numerical results show that the sensitivity of the hyperparameter estimation depends on the choice of kernels. The estimates for SE kernel are robust regardless of the prior distributions, whilst they are very different using different priors for PER kernel which implies that the prior distributions have huge impact on the estimates of the parameters. However, it is interesting to see that the GPR models always perform well in terms of predictability, despite the poor estimates of the hyperparameters in some cases. Particularly the performances of the GPR models using various priors are consistently comparable with that of the true time series model in terms of prediction. Overall, prior distributions of the hyperparameters have little impact on the performance of GPR models.

It is noted that in terms of evaluating the influences of prior distributions to the performance of GPR models, the study in this paper is far from comprehensive. A wider range of priors and kernels need to be considered, as well as more complex data, including real data. Theoretical analysis may also be of importance because it is not feasible for numerical examples to cover all scenarios.

References

References

[1] C. E. Rasmussen, Evaluation of Gaussian processes and other methods for non-linear regression, Ph.D. thesis, University of Toronto (1996).

[2] R. M. Neal, Bayesian learning for neural networks, Ph.D. thesis, University of Toronto (1995).

[3] D. MacKay, Gaussian processes-a replacement for supervised neural networks? Lecture notes for a tutorial at NIPS 1997 (1997).

[4] C. K. Williams, C. E. Rasmussen, Gaussian processes for regression, Advances in Neural Information Processing Systems.

[5] S. Brahim-Belhouari, J. Vesin, Bayesian learning using Gaussian process for time series prediction (2001).

[6] S. Brahim-Belhouari, A. Bermak, Gaussian process for nonstationary time series prediction, Computational Statistics and Data Analysis 47 (4) (2004) 705–712.

[7] A. G. Wilson, R. P. Adams, Gaussian process kernels for pattern discovery and extrapolation, arXiv preprint arXiv:1302.4245.

[8] C. E. Rasmussen, C. K. I. Williams, Gaussian Processes for Machine Learning, 12th Edition, MIT Press, Cambridge, MA, USA, 2005.

[9] R. M. Neal, Monte carlo implementation of Gaussian process models for bayesian regression and classification, arXiv preprint physics/9701026.

[10] D. J. MacKay, Introduction to Gaussian processes, NATO ASI Series F Computer and Systems Sciences 168 (1998) 133–166.

[11] D. Duvenaud, J. R. Lloyd, R. Grosse, J. B. Tenenbaum, Z. Ghahramani, Structure discovery in nonparametric regression through compositional kernel search, arXiv preprint arXiv:1302.4922.

[12] A. G. Wilson, Covariance kernels for fast automatic pattern discovery and extrapolation with Gaussian processes, Ph.D. thesis, PhD thesis, University of Cambridge (2014).

[13] S. Roberts, M. Osborne, M. Ebden, S. Reece, N. Gibson, S. Aigrain, Gaussian processes for time-series modelling, Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences 371 (1984) (2013) 20110550.

[14] P. C. Lambert, A. J. Sutton, P. R. Burton, K. R. Abrams, D. R. Jones, How vague is vague? A simulation study of the impact of the use of vague prior distributions in MCMC using WinBUGS, Statistics in Medicine 24 (15) (2005) 2401.

[15] A. Gelman, et al., Prior distributions for variance parameters in hierarchical models (comment on article by browne and draper), Bayesian Analysis 1 (3) (2006) 515–534.

[16] H. Jeffreys, The theory of probability, OUP Oxford, 1998.

[17] R. D. Rosenkrantz, ET Jaynes: Papers on probability, statistics and statistical physics, Vol. 158, Springer Science & Business Media, 2012.

[18] H. Cramér, U. Grenander, Probability and statistics: the Harald Cramér volume, Almqvist & Wiksell, 1959.